# Efficient and Modular Consensus-Free Reconfiguration for Fault-Tolerant Storage

Eduardo Alchieri, Alysson Bessani, Fabiola Greve, Joni Fraga
University of Brasilia (Brazil), University of Lisbon (Portugal),
Federal University of Bahia (Brazil), Federal University of Santa Catarina (Brazil).

arXiv:1607.05344v1 [cs.DC] 18 Jul 2016

**Abstract**—Quorum systems are useful tools for implementing consistent and available storage in the presence of failures. These systems usually comprise a static set of servers that provide a fault-tolerant read/write register accessed by a set of clients. We consider a dynamic variant of these systems and propose FREESTORE, a set of fault-tolerant protocols that emulates a register in dynamic asynchronous systems in which processes are able to join/leave the servers set during the execution. These protocols use a new abstraction called *view generators*, that captures the agreement requirements of reconfiguration and can be implemented in different system models with different properties. Particularly interesting, we present a reconfiguration protocol that is *modular, efficient, consensus-free* and *loosely coupled* with read/write protocols, improving the overall system performance.

## 1 INTRODUCTION

Quorum systems [1] are a fundamental abstraction to ensure consistency and availability of data stored in replicated servers. Apart from their use as building blocks of synchronization protocols (e.g., consensus [2], [3]), quorum-based protocols for read/write (r/w) register implementation are appealing due to their promising scalability and fault tolerance: the r/w operations do not need to be executed in all servers, but only in a subset (quorum) of them. The consistency of the stored data is ensured by the intersection between any two quorums of the system.

Quorum systems were initially studied in static environments, where servers are not allowed to join or leave the system during execution [1], [4]. This approach is not adequate for long lived systems since, given a sufficient amount of time, there might be more faulty servers than the threshold tolerated, affecting the system correctness. Beyond that, this approach does not allow a system administrator to deploy new machines (to deal with increasing workloads) or replace old ones at runtime. Moreover, these protocols can not be used in many systems where, by their very nature, the set of processes that compose the system may change during its execution (e.g., MANETs and P2P overlays).

Reconfiguration is the process of changing the set of nodes that comprise the system. Previous works proposed solutions for reconfigurable storage by implementing dynamic quorum systems [5], [6]. Those proposals rely on consensus for reconfigurations, in a way that processes agree on the set of servers (view) supporting the storage. Although adequate, since the problem of changing views resembles an agreement problem, this approach is not the most efficient or appropriate. Besides the costs of running the protocol, consensus is known to not be solvable in asynchronous environments [7]. Moreover, atomic shared memory emulation can be implemented in static asynchronous systems without requiring consensus [4].

Until recently, the distributed systems community did not know whether it would be possible to implement reconfigurations without agreement. Aguilera et al [8] answered this question affirmatively with *DynaStore*, a protocol that implements dynamic atomic storage without relying on consensus. DynaStore reconfigurations may occur at any time, generating a graph of views from which it is possible to identify a sequence of views in which clients need to execute their r/w operations (see Figure 1, left). Unfortunately, it presents two serious drawbacks: its reconfiguration and r/w protocols are strongly tied and its performance is significantly worse than consensus-based solutions in synchronous executions, which are the norm. The first issue is particularly important since it means that DynaStore r/w protocols (as well as other consensus-based works like RAMBO [5]) explicitly deal with reconfigurations and are quite different from static r/w protocols (e.g., the ABD protocol [4]).

Recently, SmartMerge [9] and SpSn [10] protocols improved DynaStore by separating the reconfiguration and r/w protocols. However, as in DynaStore, these approaches do not fully decouple de execution of r/w and reconfiguration protocols since before execute each r/w it is necessary to access an abstraction implemented by a set of static single-writer multi-reader SWMR registers to check for updates in the system. This design decision drops significantly the system performance, even in periods without reconfigurations.

In this paper we present FREESTORE, a set of algorithms for implementing fault-tolerant atomic [11] and wait-free [12] storage that allows the servers' set reconfiguration at runtime. FREESTORE is composed by two types of protocols: (1) r/w protocols and (2) the reconfiguration protocol. Read/write protocols can be adapted from static register implementations (e.g., the classical ABD [4], as used in this paper). The reconfiguration protocol – our main contribution – is used

• *This work is supported by CNPq (Brazil) through project FREESTORE (CNPq 457272/2014-7).*

to change the set of replicas supporting the storage. The key innovation of the FREESTORE reconfiguration protocol is the use of *view generators*, a new abstraction that captures the agreement requirements of reconfiguration protocols. We provide two implementations of view generators, one based on consensus and other consensus-free, and compare how efficiently they can be used to solve reconfiguration.

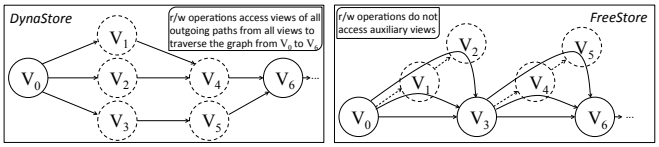
DynaStore
r/w operations access views of all outgoing paths from all views to traverse the graph from V0 to V6
r/w operations do not access auxiliary views
FreeStore

Fig. 1. View convergence strategies of DynaStore [8] (left) and FREESTORE (right). Dotted circles represent the auxiliary/non-established views that the system may experience during a reconfiguration. Solid circles represent the established/installed views process must converge to.

FREESTORE improves the state of the art in at least three aspects: *modularity*, *efficiency* and *simplicity*. The modularity of the proposed protocols is twofold: it separates the reconfiguration from the r/w protocols, allowing other static storage protocols (e.g., [13], [14]) to be adapted to our dynamic model, and introduces the notion of view generators, capturing the agreement requirements of a reconfiguration protocol. In terms of performance, both the r/w and reconfiguration protocols of FREESTORE require less communication steps than their counterparts from previous works, either consensus-based [5] or consensus-free [8], [9], [10]. In particular, FREESTORE' consensus-free reconfiguration requires less communication steps than other consensus-based reconfiguration protocols in the best case, matching the intuition that a consensus-free approach should be faster than another that rely on consensus. Finally, the consensus-free variant of FREESTORE introduces a novel reconfiguration strategy that reduces the number of intermediary views to reach an installed view with all requested updates (see Figure 1, right). This strategy is arguably easier to understand than the one used in the previous works, shedding more light on the study of consensus-free reconfiguration protocols.

In summary, the main contributions of this paper are:

1) It introduces the notion of *view generators*, an abstraction that captures the agreement requirements of a storage reconfiguration protocol and provides two implementations: the consensus-based *perfect view generator* and the consensus-free *live view generator*;
2) It shows that *safe* and *live* dynamic fault-tolerant atomic storage can be implemented using the proposed view generators and discusses the tradeoffs between them;
3) It presents FREESTORE, the first dynamic atomic storage system in which the reconfiguration protocol *(a)* can be configured to be consensus-free or consensus-based and *(b)* is fully decoupled from the r/w protocols, increasing the system performance and making it easy to adapt other static fault-tolerant register implementation to dynamic environments;
4) It presents a detailed comparison of FREESTORE with previous protocols [5], [8], [9], [10], showing that FREESTORE is faster (requires less communication steps) than these previous systems.

## 2 PRELIMINARY DEFINITIONS

### 2.1 System Model

We consider a fully connected distributed system composed by a universe of processes $U$, that can be divided in two non-overlapping subsets: an infinite set of servers $\Pi = \{1, 2, ...\}$; and an infinite set of clients $C = \{c_1, c_2, ...\}$. Clients access the storage system provided by a subset of the servers (a *view*) by executing read and write operations. Each process (client or server) of the system has a unique identifier. Servers and clients are prone to *crash failures*. Crashed processes are said to be *faulty*. A process that is not faulty is said to be *correct*. Moreover, there are *reliable channels* connecting all pairs of processes.

We assume an *asynchronous distributed system* in which there are no bounds on message transmission delays and processing times. However, each server of the system has access to a local clock used to trigger reconfigurations. These clocks are not synchronized and do not have any bounds on their drifts, being nothing more than counters that keep increasing. Besides that, there is a real-time clock not accessed by the processes, used in definitions and proofs.

### 2.2 Dynamic Storage Properties

During a dynamic system execution, a sequence of views is installed to account for replicas' joins and leaves. In the following we describe some preliminary definitions and the properties satisfied by FREESTORE.

Updates: We define $update = \{+, -\} \times \Pi$, where the tuple $\langle +, i\rangle$ (resp. $\langle -, i\rangle$) indicates that server $i$ asked to join (resp. leave) the system. A *reconfiguration* procedure takes into account updates to define a new system's configuration, which is represented by a *view*.

Views: A view $v$ is composed by a set of *updates* (represented by $v.entries$) and its associated membership (represented by $v.members$). Consequently, $v.members = \{i \in \Pi\text{: } \langle +, i\rangle \in v.entries \wedge \langle -, i\rangle \notin v.entries\}$. To simplify the notation, we sometimes use $i \in v$ and $|v|$ meaning $i \in v.members$ and $|v.members|$, respectively. Notice that a server $i$ can join and leave the system only once, but this condition can be relaxed in practice if we add an epoch number on each reconfiguration request.

We say a view $v$ is *installed* in the system if some correct server $i \in v$ considers $v$ as its *current view* and answers client r/w operations on this view. On the other hand, we say that the previous view, installed before $v$, was *uninstalled* from the system. At any time $t$, we define $V(t)$ to be the *most up-to-date view* (see definition below) installed in the system. We consider that $V(t)$ remains *active* from the time it is installed in the system until *all correct servers* of another most up-to-date view $V(t'), t' > t$, installs $V(t')$.

Comparing views: We compare two views $v_1$ and $v_2$ by comparing their $entries$. We use the notation $v_1 \subset v_2$ and $v_1 = v_2$ as an abbreviation for $v_1.entries \subset v_2.entries$ and $v_1.entries = v_2.entries$, respectively. If $v_1 \subset v_2$, then $v_2$ is *more up-to-date* than $v_1$.

Bootstrapping: We assume a non-empty initial view $V(0)$ which is known to all processes. At system startup, each server $i \in V(0)$ receives an initial view $v_0 = \{\langle +, j\rangle : j \in V(0)\}$.

Views vs. r/w operations: At any time $t$, r/w operations are executed only in $V(t)$. When a server $i$ asks to *join* the system, these operations are disabled on it until an *enable operations* event occurs. After that, $i$ remains able to process r/w operations until it asks to leave the system, which will happen when a *disable operations* event occurs.

**Definition 1 (FREESTORE properties)** FREESTORE *satisfies the following properties:*

- ***Storage Safety****: The r/w protocols satisfy the safety properties of an atomic r/w register [11].*
- ***Storage Liveness****: Every r/w operation executed by a correct client eventually completes.*
- ***Reconfiguration – Join Safety****: If a server $j$ installs a view $v$ such that $i \in v$, then server $i$ has invoked the $join$ operation or $i$ is member of the initial view.*
- ***Reconfiguration – Leave Safety****: If a server $j$ installs a view $v$ such that $i \notin v \wedge (\exists v' : i \in v' \wedge v' \subset v)$, then server $i$ has invoked the $leave$ operation.*[1]
- ***Reconfiguration – Join Liveness****: Eventually, the* enable operations *event occurs at every correct server that has invoked a $join$ operation.*
- ***Reconfiguration – Leave Liveness****: Eventually, the* disable operations *event occurs at every correct server that has invoked a $leave$ operation.*

### 2.3 Additional Assumptions for Dynamic Storage

Dynamic fault-tolerant storage protocols [5], [6], [8], [9], [10] require the following additional assumptions.

**Assumption 1 (Fault threshold)** *For each view $v$, we denote $v.f$ as the number of faults tolerated in $v$ and assume that $v.f \leq \lfloor \frac{|v.members|-1}{2} \rfloor$.*

**Assumption 2 (Quorum size)** *For each view $v$, we assume quorums of size $v.q = \lceil \frac{|v.members|+1}{2} \rceil$.*

These two previous assumptions are a direct adaptation of the optimal resilience for fault-tolerant quorum systems [4] to account for multiple views and only need to hold for views present in the generated view sequences (see Section 3).

**Assumption 3 (Gentle leaves)** *A correct server $i \in V(t)$ that asks to leave the system at time $t$ remains in the system until it knows that a more up-to-date view $V(t'), t' > t, i \notin V(t')$ is installed in the system.*

1. We can relax this property and adapt our protocol (Section 4) to allow that other process issues the leave operation on behalf of a crashed process.

The *gentle leaves* assumption ensures that a *correct* leaving server will participate in the reconfiguration protocol that installs the new view without itself. Thus, it cannot leave the system before a new view accounting for its removal is installed. Other dynamic systems require similar assumption: departing replicas need to stay available to transfer their state to arriving replicas. Notice however that the fault threshold accounts for faults while the view is being reconfigured.

**Assumption 4 (Finite reconfigurations)** *The number of updates requested in an execution is finite.*

This assumption is fundamental to guarantee that r/w operations always terminate. As in other dynamic storage systems, it ensures that a client will restart phases of an r/w due to reconfigurations a finite number of times, and thus, eventually such operation will complete. In practice, the updates proposed could be infinite, as long as each r/w is concurrent with a finite number of reconfigurations.

## 3 VIEW GENERATORS

*View generators* are distributed oracles used by servers to generate sequences of new views for system reconfiguration. This module aims to capture the agreement requirements of reconfiguration algorithms. In order to be general enough to be used for implementing consensus-free algorithms, such requirements are reflected in the sequence of generated views, and not directly on the views. This happens because, as described in previous works [5], [8], [9], [10], [15], the key issue with reconfiguration protocols is to ensure that the sequence of (possibly conflicting) views generated during a reconfiguration procedure will converge to a single view with all requested view updates.

For each view $v$, each server $i \in v$ associates a view generator $\mathcal{G}_i^v$ with $v$ in order to generate a succeeding sequence of views (possibly with a single view). Server $i$ interacts with a view generator through two primitives: (1) $\mathcal{G}_i^v.gen\_view(seq)$, called by $i$ to propose a new view sequence $seq$ to update $v$; and (2) $\mathcal{G}_i^v.new\_view(seq')$, a callback invoked by the view generator to inform $i$ that a new view sequence $seq'$ was generated for succeeding $v$. An important remark about this interface is that there is no one-to-one relationship between $gen\_view$ and $new\_view$: a server $i$ may not call the first but receive several upcalls on the latter for updating the *same view* (e.g., due to reconfigurations started by other servers). However, if $i$ calls $\mathcal{G}_i^v.gen\_view(seq)$, it will eventually receive at least one upcall to $\mathcal{G}_i^v.new\_view(seq')$.

Similarly to other distributed oracles (e.g., failure detectors [2]), view generators can implement these operations in different ways, according to the different environments they are designed to operate (e.g., synchronous or asynchronous systems). However, in this paper we consider view generators satisfying the following properties.

**Definition 2 (VIEW GENERATORS)** *A generator $\mathcal{G}_i^v$ (associated with $v$ in server $i$) satisfy the following properties:*

- ***Accuracy****: we consider two variants of this property:*

  - ***Strong Accuracy***: *for any $i, j \in v$, if $i$ receives an upcall $\mathcal{G}_i^v.new\_view(seq_i)$ and $j$ receives an upcall $\mathcal{G}_j^v.new\_view(seq_j)$, then $seq_i = seq_j$.*
  - ***Weak Accuracy***: *for any $i, j \in v$, if $i$ receives an upcall $\mathcal{G}_i^v.new\_view(seq_i)$ and $j$ receives an upcall $\mathcal{G}_j^v.new\_view(seq_j)$, then either $seq_i \subseteq seq_j$ or $seq_j \subset seq_i$.*
- ***Non-triviality***: *for any upcall $\mathcal{G}_i^v.new\_view(seq_i)$, $\forall w \in seq_i$, $v \subset w$.*
- ***Termination***: *if a correct server $i \in v$ calls $\mathcal{G}_i^v.gen\_view(seq)$, then eventually it will receive an upcall $\mathcal{G}_i^v.new\_view(seq_i)$.*

Accuracy and Non-triviality are safety properties while Termination is related with the liveness of view generation. Furthermore, the Non-triviality property ensures that generated sequences contains only updated views. Using the two variants of accuracy we can define two types of view generators: $\mathcal{P}$, the *perfect view generator*, that satisfies *Strong Accuracy* and $\mathcal{L}$, the *live view generator*, that only satisfies *Weak Accuracy*. Our implementation of $\mathcal{P}$ requires consensus, while $\mathcal{L}$ can be implemented without such strong synchronization primitive.

### 3.1 Perfect View Generators – $\mathcal{P}$

*Perfect view generators* ensure that a single sequence of views will be eventually generated at all servers of $v$ for updating such view. Our implementation for $\mathcal{P}$ (Algorithm 1) uses a deterministic Paxos-like consensus protocol [3] that assumes a partially synchronous system model [16]. The algorithm is straightforward: when $\mathcal{G}_i^v.gen\_view(seq)$ is called by server $i$, it first verifies if $\forall w \in seq : v \subset w$ and, if it is the case, propose $seq$ as its value for the consensus reconfiguring $v$. This verification aims to ensure that Non-triviality is respected, i.e., only updated views will be proposed. When a value $seq'$ is decided in the consensus associated with view $v$, an upcall to $\mathcal{G}_i^v.new\_view(seq')$ is executed. The Termination and Strong Accuracy properties comes directly from the Agreement and Termination properties of the underlying consensus algorithm [2], [3].

**Algorithm 1** $\mathcal{P}$ associated with $v$ - server $i \in v$.

**procedure** $\mathcal{G}_i^v.gen\_view(seq)$
1) **if** $\forall w \in seq : v \subset w$ **then** $Paxos^v.propose(seq)$
**upon** $Paxos^v.learn(seq')$
2) $\mathcal{G}_i^v.new\_view(seq')$

Correctness: The following theorem proves that Algorithm 1 satisfies the properties of $\mathcal{P}$.

**Theorem 1** *Algorithm 1 implements a perfect view generator in a partially synchronous system model.*

*Proof: Termination* and *Strong Accuracy* are ensured by classical consensus properties, namely, Termination and Agreement, respectively. *Non-Triviality* is ensured by the verification of line 1 and the consensus Validity property (the decided value is one of the proposals). □

### 3.2 Live View Generators – $\mathcal{L}$

Algorithm 2 presents an implementation for *Live view generators* ($\mathcal{L}$). Our algorithm does not require a consensus building block, being thus implementable in asynchronous systems. On the other hand, it can generate different sequences in different servers for updating the same view. We bound such divergence by exploiting Assumptions 1 and 2, which ensure that any quorum of the system will intersect in at least one correct server, making any generated sequence for updating $v$ be contained in any other posterior sequence generated for $v$ (Weak Accuracy). Furthermore, the servers keep updating their proposals until a quorum of them converges to a sequence containing all proposed views (or combinations of them), possibly generating some intermediate sequences before this convergence.

To generate a new sequence of views, a server $i \in v$ uses an auxiliary function $most\_updated$ to get the most up-to-date view in a sequence of views (i.e., the view that is not contained in any other view of the sequence). Moreover, each server keeps two local variables: $\texttt{SEQ}^v$ – the last view sequence proposed by the server – and $\texttt{LCSEQ}^v$ – the last sequence this server converged.

When server $i \in v$ starts its view generator, it first verifies (1) if it already made a proposal for updating $v$ and (2) if the sequence being proposed contains only updated views (line 1). If these two conditions are met, it sends its proposal to the servers in $v$ (lines 2-3).

**Algorithm 2** $\mathcal{L}$ associated with $v$ - server $i \in v$.

**functions:** Auxiliary functions
  $most\_updated(seq) \equiv w: (w \in seq) \wedge (\nexists w' \in seq: w \subset w')$
**variables:** Sets used in the protocol
  $\texttt{SEQ}^v \leftarrow \emptyset$ // proposed view sequence
  $\texttt{LCSEQ}^v \leftarrow \emptyset$ // last converged view sequence known
**procedure** $\mathcal{G}_i^v.gen\_view(seq)$
1) **if** $\texttt{SEQ}^v = \emptyset \wedge \forall w \in seq : v \subset w$ **then**
2) $\texttt{SEQ}^v \leftarrow seq$
3) $\forall j \in v, send\langle$SEQ-VIEW, $\texttt{SEQ}^v\rangle$ to $j$
**upon receipt of** $\langle$SEQ-VIEW, $seq\rangle$ from $j$
4) **if** $\exists w \in seq: w \notin \texttt{SEQ}^v$ **then**
5) **if** $\exists w, w': w \in seq \wedge w' \in \texttt{SEQ}^v \wedge w \not\subset w' \wedge w' \not\subset w$ **then**
6) $w \leftarrow most\_updated(\texttt{SEQ}^v)$
7) $w' \leftarrow most\_updated(seq)$
8) $\texttt{SEQ}^v \leftarrow \texttt{LCSEQ}^v \cup \{w.entries \cup w'.entries\}$
9) **else**
10) $\texttt{SEQ}^v \leftarrow \texttt{SEQ}^v \cup seq$
11) $\forall k \in v, send\langle$SEQ-VIEW, $\texttt{SEQ}^v\rangle$ to $k$
**upon receipt of** $\langle$SEQ-VIEW, $\texttt{SEQ}^v\rangle$ from $v.q$ servers in $v$
12) $\texttt{LCSEQ}^v \leftarrow \texttt{SEQ}^v$
13) $\forall k \in v, send\langle$SEQ-CONV, $\texttt{SEQ}^v\rangle$ to $k$
**upon the receipt of** $\langle$SEQ-CONV, $seq'\rangle$ from $v.q$ servers in $v$
14) $\mathcal{G}_i^v.new\_view(seq')$

Different servers of $v$ may propose sequences containing different views and therefore these views need to be organized in a sequence. When a server $i \in v$ receives a proposal for a view sequence from $j \in v$, it verifies if this proposal contains some view it did not know yet (line 4).[2] If this is the case, $i$ updates its proposal ($\texttt{SEQ}^v$) according to two mutually exclusive cases:

- CASE 1 [There are *conflicting views* in the sequence proposed by $i$ and the sequence received from $j$ (lines

2. Notice that it might happen that a server receives this message even if its view generator was not yet initialized, i.e., $\texttt{SEQ}^v = \emptyset$.

5-8)]: In this case $i$ creates a new sequence containing the last converged sequence ($\texttt{LCSEQ}^v$) and a new view with the union of the two most up-to-date conflicting views. This maintains the containment relationship between any two generated view sequences.

- CASE 2 [The sequence proposed by $i$ and the received sequence *can be composed in a new sequence* (lines 9-10)]: In this case the new sequence is the union of the two known sequences.

In both cases, a new proposal containing the new sequence is disseminated (line 11). When $i$ receives the same proposal from a quorum of servers in $v$, it *converges* to $\texttt{SEQ}^v$ and stores it in $\texttt{LCSEQ}^v$, informing other servers of $v$ about it (lines 12-13). When $i$ knows that a quorum of servers of $v$ converged to some sequence $seq'$, $i$ *generates* $seq'$ (line 14).

Correctness: The algorithm ensures that if a quorum of servers converged to a sequence $seq'$ (lines 12-13), then *(1)* such sequence will be generated (line 14) and *(2)* any posterior sequence generated will contain $seq'$ (lines 8 and 10), ensuring Weak Accuracy. This holds due to the quorum intersection: at least one correct server needs to participate in the generation of both sequences and this server applies the rules in Cases 1 and 2 to ensure that sequences satisfy the containment relationship. The servers in $v$ generate at most $|v| - v.q + 1$ sequences of views, such that $seq_1 \subset ... \subset seq_{|v|-v.q+1}$. The sequence $seq_1$ is obtained by merging the proposals of a quorum of servers; the other sequences may be built using $seq_1$ plus different proposals from each of the $|v| - v.q$ servers that did not participate in the generation of $seq_1$. The Termination property is ensured by the fact that *(1)* each server makes at most one initial proposal (lines 1-3); *(2)* servers keep updating their proposals until a quorum agree on some proposal; and *(3)* there is always a quorum of correct servers in $v$.

The following lemmata and theorem prove that Algorithm 2 satisfies the properties of $\mathcal{L}$.

**Lemma 1** *(Weak Accuracy) Consider the view generators $\mathcal{G}^v$ implemented by Algorithm 2 associated with view $v$. For any $i, j \in v$, if $i$ receives an upcall $\mathcal{G}_i^v.new\_view(seq_i)$ and $j$ receives an upcall $\mathcal{G}_j^v.new\_view(seq_j)$, then either $seq_i \subseteq seq_j$ or $seq_j \subset seq_i$.*

*Proof:* Consider that $\mathcal{G}^v$ generates sequences $seq_i$ and $seq_j$ at servers $i, j \in v$ by calling $\mathcal{G}_i^v.new\_view(seq_i)$ and $\mathcal{G}_j^v.new\_view(seq_j)$, respectively. If $\mathcal{G}_i^v$ generated $seq_i$, there is a quorum of servers in $v$ that sent SEQ-VIEW messages with $seq_i$ (lines 12-13). The same can be said about the generation of $seq_j$ by $\mathcal{G}_j^v$: a quorum of servers in $v$ proposed $seq_j$ in SEQ-VIEW messages. Due to the intersection between quorums, there will be at least one correct server $k$ that sent SEQ-VIEW messages both for $seq_i$ and $seq_j$. Consider that (1) server $k$ first sent $seq_i$, (2) $i$ receives a quorum of these proposals and converges to $seq_i$ and (3) $j$ made an initial proposal $seq_j^p$. Then, we have three cases:

1) Case $seq_i = seq_j^p$, we have two cases: *(i)* $j$ receives the same messages that made $i$ converge to $seq_i$ and $\mathcal{G}_j^v$ generates $seq_j = seq_i = seq_j^p$ (lines 12-14); or *(ii)* $j$ receives a proposal containing other views that are not in $seq_j^p$ and updates its proposal before the generation of $seq_j$ (lines 8 and 10). In *(ii)*, since $seq_i$ was generated, at least a quorum of servers in $v$ store $seq_i$ in their $\texttt{LCSEQ}^v$ (line 12), which means $seq_j$ must contain $seq_i$.
2) Case $seq_i \neq seq_j^p$ and *conflicting*, the reception of SEQ-VIEW from $k$ makes $j$ proposes a new sequence containing its $\texttt{LCSEQ}^v$ ($lcseq_j$) plus a new view $w$, representing the union of all update requests (line 8): $lcseq_j \rightarrow w$. Once $seq_i$ was generated, at least a quorum of servers in $v$ store $seq_i$ in their $\texttt{LCSEQ}^v$ (line 12). Since $\texttt{LCSEQ}^v$ is updated only if SEQ-VIEW messages with the same sequence are received from a quorum of servers (line 12), it will not contain conflicting sequences in different servers. Consequently, the generated sequence $seq_j$ must contain $seq_i$.
3) Case $seq_i \neq seq_j^p$ and *not conflicting*, $j$ updates its proposal to a sequence $seq_j$ representing the union of $seq_i$ and $seq_j^p$ (line 10). After that, $j$ sends a SEQ-VIEW message with $seq_j$ to other servers that eventually update their proposals to this sequence. Consequently, $\mathcal{G}_j^v$ generated a $seq_j$ such that $seq_i \subset seq_j$.

Therefore, in the three possible cases $seq_i \subseteq seq_j$. Using similar arguments it is possible to prove the symmetrical execution: if server $k$ sent a SEQ-VIEW message first for $seq_j$ and then for $seq_i$, we have that $seq_j \subseteq seq_i$. Consequently, we have that either $seq_i \subseteq seq_j$ or $seq_j \subset seq_i$. □

**Lemma 2** *(Termination) Consider the view generators $\mathcal{G}^v$ implemented by Algorithm 2 associated with view $v$. If a correct server $i \in v$ calls $\mathcal{G}_i^v.gen\_view(seq)$, then eventually it will receive an upcall $\mathcal{G}_i^v.new\_view(seq_i)$.*

*Proof:* When $\mathcal{G}_i^v.gen\_view(seq)$ is called, server $i$ sends a message with its proposal $seq$ (line 3). Since the channels are reliable, messages sent by correct processes are received by correct processes (lines 4-11). For $\mathcal{G}_i^v$ to generate a new sequence of views $seq_i$ (line 14), at least $v.q$ servers should converge to $seq_i$ (lines 12-13). During this convergence (lines 4-11), $i$ changes its proposed sequence ($\texttt{SEQ}^v$) with proposals received from other servers (line 11). Since there are at least $v.q$ correct servers in $v$ and each server makes a single initial proposal (line 1), at least $v.q$ servers are going to converge to the same sequence. Thus, $i$ will eventually generate $seq_i$ by executing $\mathcal{G}_i^v.new\_view(seq_i)$ (line 14). □

**Lemma 3** *(Bounded Sequences for $\mathcal{L}$) Consider the view generators $\mathcal{G}^v$ implemented by Algorithm 2 associated with view $v$. The number of different view sequences generated by $\mathcal{G}^v$ is bounded.*

*Proof:* $\mathcal{G}^v$ must collect a quorum of proposals for a sequence $seq$ to generate it (lines 12-14). Each server $i \in v$ makes a single initial proposal (line 1) and updates it only by processing proposals from other server of $v$ (lines 8 and 10). Consequently, $\mathcal{G}^v$ must process the proposals from at least a quorum of servers to generate $seq$, being $|v| - v.q$ the number of servers whose proposals have not been processed to generate $seq$. In the worst case, each of these servers make a different proposal and such proposal is combined with the last

generated sequence to generate a new sequence. Consequently, $\mathcal{G}^v$ generates at most $|v| - v.q + 1$ sequences. □

**Theorem 2** *Algorithm 2 implements a live view generator.*

*Proof: Weak Accuracy* and *Termination* follows directly from Lemmata 1 and 2, respectively. The *Non-Triviality* property is ensured by the verification of line 1 and the way further proposals are defined (lines 8 and 10). □

# 4 FREESTORE RECONFIGURATION

A server running the FREESTORE reconfiguration algorithm uses a view generator associated with its current view $cv$ to process one or more reconfiguration requests (joins and leaves) that will lead the system from $cv$ to a new view $w$. Algorithm 3 describes how a server $i$ executes reconfigurations. In the following sections we first describe how view generators are started and then we proceed to discuss the behavior of this algorithm when started with either $\mathcal{L}$ (Section 4.2) or $\mathcal{P}$ (Section 4.3).

## 4.1 View Generator Initialization

Algorithm 3 (lines 1-7) describes how a server $i$ processes reconfiguration requests and starts a view generator associated with its current view $cv$. A server $j$ that wants to join the system needs first to find the current view $cv$ and then to execute the $join$ operation (lines 1-2), sending a tuple $\langle +, j \rangle$ to the members of $cv$. Servers leaving the system do something similar, through the $leave$ operation (lines 3-4). When $i$ receives a reconfiguration request from $j$, it first verifies if the requesting server is using the same view as itself; if this is not the case, $i$ replies its current view to $j$ (omitted from the algorithm for brevity). If they are using the same view and $i$ did not execute the requested reconfiguration before, it stores the reconfiguration tuple in its set of pending updates `RECV` and sends an acknowledgment to $j$ (lines 5-6).

For the sake of performance, a local *timer* has been defined in order to periodically process the updates requested in a view, that is, the next system reconfiguration. A server $i \in cv$ starts a reconfiguration for $cv$ when its timer expires and $i$ has some pending reconfiguration requests (otherwise, the timer is renewed). The view generator is started with a sequence containing a single view representing the current view plus the pending updates (line 7).

## 4.2 Reconfiguration using $\mathcal{L}$

This section presents the reconfiguration processing when Algorithm 3 is instantiated with *live view generators* that may generate different views at different servers. In this case, the algorithm ensures that the system *installs* an unique sequence of views.

Overview: Given a sequence $seq : v_1 \rightarrow \ldots \rightarrow v_k \rightarrow w$ generated for updating a view $v$, only the last view $w$ will be installed in the system. The other $k$ *auxiliary views* are used only as intermediate steps for installing $w$. The use of auxiliary views is fundamental to ensure that no write operation executed in any of the views of $seq$ (in case they are installed in some server) "is lost" by the reconfiguration processing. This is done through the execution of a *"chain of reads"* in which the servers of $v$ transfer their state to the servers of $v_1$, which transfer their state to the servers of $v_2$ and so on until the servers of $w$ have the most up-to-date state. To avoid consistency problems, r/w operations are disabled during each of these state transfers.

It is important to remark that since we do not use consensus (nor solve an equivalent problem), a subsequence $seq' : v_1 \rightarrow \ldots \rightarrow v_j, j \leq k$, of $seq$ may lead to the installation of $v_j$ in some servers that did not know $seq$ and that these servers may execute r/w operations in this view. However, the algorithm ensures these servers eventually will reconfigure from $v_j$ to the most up-to-date view $w$.

The protocol: Algorithm 3 (lines 8-28) presents the core of the FREESTORE reconfiguration protocol. This algorithm uses an auxiliary function $least\_updated$ to obtain the least updated view in a sequence of views (i.e., the one that is contained in all other views of the sequence) and two local variables: the aforementioned `RECV` – used to store pending reconfiguration requests – and $cv$ – the current view of the server (initially $v_0$).

When the view generator associated with some view $ov$ reports[3] the generation of a sequence of views $seq$, the server obtains the least updated view $w$ of $seq$ and proposes this sequence for updating $ov$ through an INSTALL-SEQ message sent to the servers of both, $ov$ and $w$. To ensure all correct servers in this group will process the reconfiguration, we employ a *reliable multicast* primitive [17] (lines 8-9), that can be efficiently implemented in asynchronous systems with a message retransmission at the receivers before its delivery (i.e., in one communication step) [2], [17].

The current view is updated through the execution of lines 10-28. First, if the server is a member of the view being updated $ov$, it must send its state (usually, the register's value and timestamp) to the servers in the new view to be installed (lines 10-12). However, if the server will be updating its current view (i.e., if $w$ is more up-to-date than $cv$) it first needs to stop executing client r/w operations (line 11) and enqueue these operations to be executed when the most up-to-date view in the sequence is installed (line 25, as discussed bellow). A server will update its current view only if the least updated view $w$ of the proposed sequence is more up-to-date than its current view $cv$ (line 13). If this is the case, either (1) server $i$ will be in the next view (lines 14-25) or (2) not (line 26-28).

- CASE 1 [$i$ will be in the next view (it may be joining the system)]: If the server will be in the next view $w$, it first waits for the state from a quorum of servers from the previous view $ov$ and then defines the current value and timestamp of the register (lines 15-16), similarly to what is done in the 1st phase of a read operation (see Section 5). After ensuring that its state is updated, the server updates $cv$ to $w$ and, if it is joining the system, it enables the processing of r/w operations (which will be queued until line 25 is executed). Furthermore, the server

3. We use $ov$ instead of $cv$ because view generators associated with *old views* ($ov \subseteq cv$) still active can generate new sequences.

**Algorithm 3** FREESTORE reconfiguration - server $i$.

**functions:** Auxiliary functions
$least\_updated(seq) \equiv w{:}\ (w \in seq) \wedge (\nexists w' \in seq{:}\ w' \subset w)$
**variables:** Sets used in the protocol
$cv \leftarrow v_0$ // the system current view known by $i$
RECV $\leftarrow \emptyset$ // set of received updates
**procedure** $join()$
1) $\forall j \in cv, send\langle$RECONFIG, $\langle +, i\rangle, cv\rangle$ to $j$
2) **wait** for $\langle$REC-CONFIRM$\rangle$ replies from $cv.q$ servers in $cv$
**procedure** $leave()$
3) $\forall j \in cv, send\langle$RECONFIG, $\langle -, i\rangle, cv\rangle$ to $j$
4) **wait** for $\langle$REC-CONFIRM$\rangle$ replies from $cv.q$ servers in $cv$
**upon receipt of** $\langle$RECONFIG, $\langle *, j\rangle, cv\rangle$ from $j$ and $\langle *, j\rangle \notin cv$
5) RECV $\leftarrow$ RECV $\cup \{\langle *, j\rangle\}$
6) $send\langle$REC-CONFIRM$\rangle$ to $j$
**upon** ($timeout$ for $cv$) $\wedge$(RECV $\neq \emptyset$)
7) $\mathcal{G}_i^{cv}.gen\_view(\{cv \cup$ RECV$\})$
**upon** $\mathcal{G}_i^{ov}.new\_view(seq)$ //$\mathcal{G}$ generates a new sequence of views to update $ov$ (usually $ov = cv$)
8) $w \leftarrow least\_updated(seq)$ //the next view in the sequence $seq$
9) *R-multicast*($\{j{:}\ j \in ov \vee j \in w\}$,$\langle$INSTALL-SEQ, $w, seq, ov\rangle$)
**upon** *R-delivery*($\{j{:}\ j \in ov \vee j \in w\}$,$\langle$INSTALL-SEQ, $w, seq, ov\rangle$)
10) **if** $i \in ov$ **then** //$i$ is member of the previous view in the sequence
11) **if** $cv \subset w$ **then** *stop the execution of r/w operations* //if $w$ is more up-to-date than $cv$, enqueue r/w operations while the state transfer is done
12) $\forall j \in w, send\langle$STATE-UPDATE, $\langle val, ts\rangle$, RECV$\rangle$ to $j$ //$i$ sends its state to servers in the next view
13) **if** $cv \subset w$ **then** //$w$ is more up-to-date than $cv$ and the system will be reconfigured from $cv$ to $w$
14) **if** $i \in w$ **then** //$i$ is in the new view
15) **wait** for $\langle$STATE-UPDATE, $*, *\rangle$ messages from $ov.q$ servers in $ov$
16) $\langle val, ts\rangle \leftarrow \langle val_h, ts_h\rangle$, pair with highest timestamp among the ones received //$i$ updates its state...
17) RECV $\leftarrow$ RECV $\cup$ {update requests from STATE-UPDATE messages} $\setminus\ w.entries$
18) $cv \leftarrow w$ //... and updates its current view to $w$
19) **if** $i \notin ov$ **then** *enable operations* //$i$ joining the system
20) $\forall j \in ov \setminus cv, send\langle$VIEW-UPDATED, $cv\rangle$ to $j$ //inform servers in $ov \setminus cv$ that they can leave
21) **if** $(\exists w' \in seq{:}\ cv \subset w')$ **then** //there are views more up-to-date than $cv$ in $seq$...
22) $seq' \leftarrow \{w' \in seq{:}\ cv \subset w'\}$ //... gather these views...
23) $\mathcal{G}_i^{cv}.gen\_view(seq')$ //... and propose them (going back to Algorithm 2)
24) **else**
25) *resume the execution of r/w operations* in $cv = w$ and start a timer for $cv$ //$w$ **is installed**
26) **else** //$i$ is leaving the system
27) *disable operations*
28) **wait** for $\langle$VIEW-UPDATED, $w\rangle$ messages from $w.q$ servers in $w$ and then **halt**

informs leaving servers that its current view was updated (line 20). The final step of the reconfiguration procedure is the verification if the new view will be installed or not (i.e., an auxiliary view). If $cv = w$ is not the most up-to-date view of the generated sequence $seq$, a new sequence with the views more up-to-date than $cv$ will be proposed for updating it (lines 21-23). Otherwise, $cv$ is installed and server $i$ resumes processing r/w operations (lines 24-25).

- CASE 2 [$i$ is leaving the system]: A server that is leaving the system only halts after ensuring that the view $w$ to which it sent its state was started in a quorum of servers (lines 27-28).

Although the algorithm (in some sense) restarts itself in line 23, the fact that the number of reconfiguration requests is finite (Assumption 4) makes the reconfiguration procedure eventually terminate. Furthermore, since all sequences generated by $\mathcal{L}$ are compatible, when the reconfiguration terminates in all correct servers, they will have installed the same view with all requested updates.

Optimization: It is possible to reduce one communication step in the FREESTORE reconfiguration protocol (Algorithm 3) when using $\mathcal{L}$ (Algorithm 2). The idea is to use the SEQ_CONV messages of $\mathcal{L}$ to inform the servers about a new sequence to be installed, making thus the INSTALL_SEQ message unnecessary.

An Execution Example: Figure 2 illustrates an execution of FREESTORE reconfiguration with $\mathcal{L}$ that makes a set of servers converge to a single view, even in the presence of conflicting proposals. Let $v_0 = \{1, 2, 3\}$ be the initial view of a system using our algorithm with the $\mathcal{L}$ view generator. In this example, ① servers $1, 2, 3$ receive the $join$ request from process $4$ before they start their view generators, while the join request from server $5$ was received only by server $1$. Consequently, the view generators associated with $v_0$ ($L(v_0)$) are started at servers $2, 3$ with a sequence containing a view $v_1 = \{1, 2, 3, 4\}$ and at server $1$ with a sequence containing a view $v_2 = \{1, 2, 3, 4, 5\}$. The Weak Accuracy of $L(v_0)$ allows ② servers $2, 3$ to generate a sequence $seq_1 : v_1$ and later servers $1, 2, 3$ may generate the sequence $seq_2 : v_1 \rightarrow v_2$ also for updating $v_0$.

In this execution, ② servers $3, 4$ receive INSTALL-SEQ messages for $seq_1$ and ③ install $v_1$ before the reception of INSTALL-SEQ messages for $seq_2$. Additionally, ③ servers $3, 4$ start their view generators $L(v_1)$ for updating $v_1$ to add server 6, which sends a join request to them. More specifically, they propose a sequence $seq_3 : v_3$ with $v_3 = \{1, 2, 3, 4, 6\}$ (line 7). Concurrently, servers $1, 2$ receive the INSTALL-SEQ messages for $seq_2$ before receiving these messages for $seq_1$.

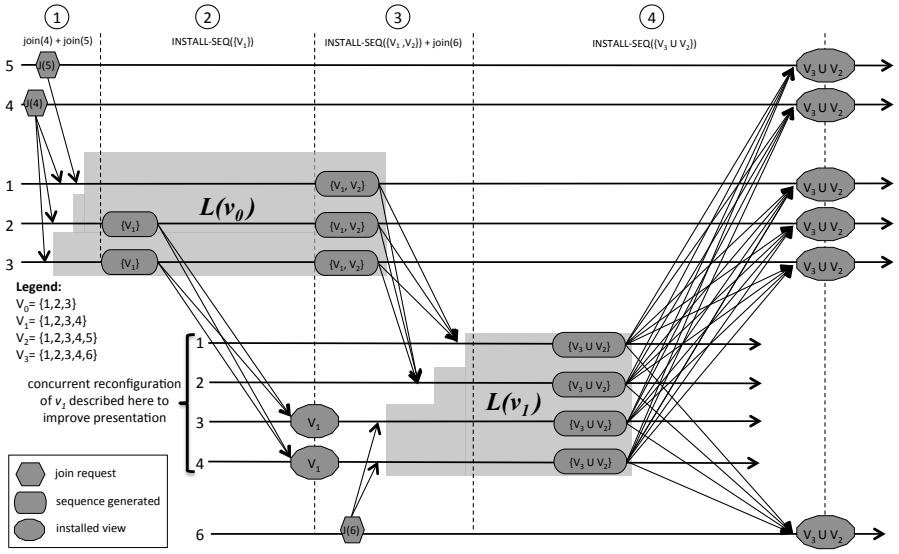


Fig. 2. Processes always converge to a single view, even in the presence of conflicting proposals.

Consequently, ③ $v_1$ will be an auxiliary view at servers $1, 2$ and they will start $L(v_1)$ with a sequence containing $seq_4 : v_2$ (line 23).

In this scenario, servers $1, 2, 3, 4$ (from $v_1$) have to handle conflicting sequence proposals for updating $v_1$, since neither $v_3 \subset v_2$ nor $v_2 \subset v_3$. Depending on the schedule of messages received during the convergence of $L(v_1)$ (i.e., the composition of the first quorum of SEQ-VIEW messages received by processes in $v_1$), the new sequence could be either (1) $v_2$, (2) $v_3$, (3) $v_2 \rightarrow (v_3 \cup v_2)$, (4) $v_3 \rightarrow (v_3 \cup v_2)$ or (5) $v_3 \cup v_2$. However, the Weak Accuracy property of $L(v_1)$ ensures that servers will not install both $v_3$ and $v_2$ (the conflicted views). Since neither $\{1, 2\}$ nor $\{3, 4\}$ form a quorum of $v_1$ (which requires three servers), they must deal with the conflicting proposals before converging to some sequence (lines 4-8 of Algorithm 2). Consequently, the only allowed sequence is (5). Thus, ④ $L(v_1)$ generates a sequence $seq_5 : v_3 \cup v_2$ and all servers install the view $v_3 \cup v_2 = \{1, 2, 3, 4, 5, 6\}$.

Correctness (full proof in Section 6.1): Algorithm 3 ensures that *an unique sequence of views is installed in the system* by the following: (1) if a view $w$ is installed, any previously installed view $w' \subset w$ is uninstalled and will not be installed anymore (lines 11, 18 and 25); (2) consequently, no view more up-to-date than $w$ is installed and, thus, the installed views form an unique sequence. Moreover, by Assumption 4 the reconfiguration procedure always terminate by installing a final view $v_{final}$. The *Storage Safety* and *Storage Liveness* properties are discussed in Section 5. The remaining properties of Definition 1 are ensured as follows. *Reconfiguration Join/Leave Safety* are trivially ensured by the fact that only a server $i$ sends the reconfiguration request $\langle +, i\rangle/\langle -, i\rangle$ (lines 1-4). *Reconfiguration Join/Leave Liveness* are ensured by the fact that if an update request from a server $i$ is stored in RECV of a quorum, then it will be processed in the next reconfiguration. This happens because a quorum with the same proposal is required to generate some view sequence (Algorithm 2), and this quorum intersects with the servers that acknowledged request of $i$ in at least one correct server. Moreover, update requests received during a reconfiguration are sent to the next view (lines 12-17).

### 4.3 Reconfiguration using $\mathcal{P}$

If $\mathcal{P}$ is used with the FREESTORE reconfiguration protocol (Algorithm 3), all generators will generate the same sequence of views (Strong Accuracy) with a single view $w$ (obtained in line 8). This will lead the system directly from its current view $cv$ to $w$ (lines 22-23 will never be executed).

## 5 READ AND WRITE PROTOCOLS

This section discusses how a static storage protocol can be adapted to dynamic systems by using FREESTORE reconfigurations. Since reconfigurations are decoupled from r/w protocols, they are very similar to their static versions. In a nutshell, there are two main requirements for using our reconfiguration protocol. First, each process (client or server) needs to handle a current view variable $cv$ that stores the most up-to-date view it knows. All r/w protocol messages need to carry $cv$ and clients update it as soon as they discover that there is a more recent view installed in the system. The servers, on the other hand, reject any operation issued to an old view, and reply their current view to the issuing client, which updates its $cv$. The client restarts the phase of the operation it is executing if it receives a view different from the one that it handles.

The second requirement is that, before accessing the system, a client must obtain the system's current view. This can be done by making servers put the current view in a directory service [15] or making the client flood the network asking for it. Notice this is an intrinsic problem for any dynamic system and similar assumptions are required in previous reconfiguration protocols [5], [6], [8], [9], [10].

In this paper we extend the classical ABD algorithm [4] for supporting multiple writers and to work with the FREESTORE

**Algorithm 4** Algorithm executed at client $c$

```
procedure write(value)
    Phase 1
    ABD Phase 1
 1) send⟨READ_TS⟩ to each i ∈ cv
 2) repeat
 3)     wait reply ⟨READ_TS_REP, ts_i, v_i⟩ from i ∈ cv
 4)     RTS_i ← ⟨READ_TS_REP, ts_i, v_i⟩
 5) until |RTS| ≥ cv.q
 6) ts ← succ(max(RTS), c)
    FreeStore Phase 1 (updated view verification)
 7) if ∃⟨READ_TS_REP, ts_j, v_j⟩ ∈ RTS: cv ≠ v_j then
 8)     if cv ⊂ v_j then cv ← v_j
 9)     restart phase 1

    Phase 2
    ABD Phase 2
10) send⟨WRITE, value, ts⟩ to each i ∈ cv
11) repeat
12)     wait reply ⟨WRITE_REP, v_i⟩ from i ∈ cv
13)     WROTE_i ← ⟨WRITE_REP, v_i⟩
14) until |WROTE| ≥ cv.q
    FreeStore Phase 2 (updated view verification)
15) if ∃⟨WRITE_REP, v_j⟩ ∈ WROTE: cv ≠ v_j then
16)     if cv ⊂ v_j then cv ← v_j
17)     restart phase 2
```

```
procedure read()
    Phase 1
    ABD Phase 1
18) send⟨READ⟩ to each i ∈ cv
19) repeat
20)     wait reply ⟨READ_REP, ⟨val, ts⟩_i, cv_i⟩ from i ∈ cv
21)     READ_i ← ⟨READ_REP, ⟨val, ts⟩_i, cv_i⟩
22) until |READ| ≥ cv.q
23) ⟨val_h, ts_h⟩ ← maxTS(READ)
    FreeStore Phase 1 (updated view verification)
24) if ∃⟨READ_REP, ⟨val, ts⟩_j, cv_j⟩ ∈ READ: cv ≠ v_j then
25)     if cv ⊂ v_j then cv ← v_j
26)     restart phase 1

    Phase 2 (write-back phase)
    ABD Phase 2
27) if ∃⟨READ_REP, ⟨∗, ts⟩_∗, ∗⟩ ∈ READ: ts ≠ ts_h then
28)     send⟨WRITE, val_h, ts_h⟩ to each i ∈ cv
29)     repeat
30)         wait reply ⟨WRITE_REP, cv_i⟩ from i ∈ cv
31)         WROTE_i ← ⟨WRITE_REP, cv_i⟩
32)     until |WROTE| ≥ cv.q
    FreeStore Phase 2 (updated view verification)
33) if ∃⟨WRITE_REP, v_j⟩ ∈ WROTE: cv ≠ v_j then
34)     if cv ⊂ v_j then cv ← v_j
35)     restart phase 2
36) return val_h
```

**Algorithm 5** Algorithm executed at server $i$.

```
    ABD server algorithm
    Upon receipt of ⟨READ_TS⟩ from client c
 1) send⟨READ_TS_REP, ts, cv⟩ to c
    Upon receipt of ⟨WRITE, val_w, ts_w⟩ from client c
 2) if ts_w > ts then
 3)     val ← val_w
 4)     ts ← ts_w
 5) send⟨WRITE_REP, cv⟩ to c
    Upon receipt of ⟨READ⟩ from client c
 6) send⟨READ_REP, ⟨val, ts⟩, cv⟩ to c
```

reconfiguration. In the following we highlight the main aspects of these r/w protocols (algorithms 4 and 5). In the algorithms, we highlight the ABD code as well as the code related with FREESTORE.

The protocols to read and write from the dynamic distributed storage work in phases. Each phase corresponds to an access to a quorum of servers in $cv$. The *read protocol* works as follows:

- (1ST PHASE) a reader client requests a set of tuples $\langle val, ts\rangle$ from a quorum of servers in $cv$ ($val$ is the value that the server stores and $ts$ is its associated timestamp) and selects the one with highest timestamp $\langle val_h, ts_h\rangle$; the operation ends and returns $val_h$ if all returned pairs are equal, which happens in executions without write contention or failures;
- (2ND PHASE) otherwise, the reader client performs an additional *write-back phase* in the system and waits for confirmations from a quorum of servers in $cv$ before returning $val_h$.

The *write protocol* works in a similar way:

- (1ST PHASE) a writer client obtains a set of timestamps from a quorum of servers in $cv$ and chooses the highest, $ts_h$; the timestamp to be written $ts$ is defined by incrementing $ts_h$ and concatenating the writer id in its lowest bits;
- (2ND PHASE) the writer sends a tuple $\langle val, ts\rangle$ to the servers of $cv$, writing $val$ with timestamp $ts$, and waits for confirmations from a quorum.

The proposed decoupling of r/w protocols from the reconfiguration makes it easy to adapt other static fault-tolerant register implementation to dynamic environments and makes it possible to run r/w operations in parallel with reconfigurations. This happens because, differently from previous approaches [5], [8], [9], [10], where r/w operations may be executed in multiple views of the system, in FREESTORE a client executes these operations only in the most up-to-date view installed in the system, which is received directly from the servers. To enable this, it is required that r/w operations to be blocked during the state transfer between views (Algorithm 3).

Correctness (full proof in Section 6.2): The above algorithms implement atomic storage in the absence of reconfigurations [4]. When integrated with FREESTORE, they also have to satisfy the *Storage Safety* and *Storage Liveness* properties of Definition 1. We start by discussing the *Storage Liveness* property, which follows directly from the termination of reconfigurations. As discussed before, clients' r/w operations are concluded only if they access a quorum of servers using the same view as the client, otherwise they are restarted. By Assumption 4, the system reconfiguration always terminate by installing some final view. Consequently, a client will restart phases of its operations a finite number of times until this final view is installed in a quorum. *Storage Safety* comes directly from three facts: (1) a r/w operation can only be executed in a single *installed* view (i.e., all servers in the quorum of the operation have the same installed current view), (2) all installed views form a unique sequence, and (3) any operation

executed in a view $v$ will still be "in effect" when a more up-to-date view $w$ is installed. More precisely, assume $\langle val, ts \rangle$ is the last value read or written in $v$, and thus it was stored in $v.q$ servers from $v$. During a reconfiguration from $v$ to $w$, r/w operations are disabled until all servers of $v$ send $\langle val, ts \rangle$ to the servers of $w$ (line 12), which terminate the reconfiguration only after receiving the state from a quorum of servers of $v$. Consequently, all servers who reconfigure to $w$ will have $\langle val, ts \rangle$ as its register's state (lines 15-16). This ensures that any operation executed in $w$ will not miss the operations executed in $v$.

# 6 FREESTORE CORRECTNESS PROOFS

## 6.1 Basic Properties

The following lemmata show a sketch of the proofs that Algorithm 3 with $\mathcal{L}$ installs an unique sequence of views in the systems and that the reconfiguration procedure always terminate.[4]

Recall that a view $v$ is *installed* in the system if some correct server $i \in v$ considers $v$ as its *current view* (var $cv$ in line 18 - Algorithm 3) and enable r/w operations in this view (line 19). When this happens, we say that the previous view, which was installed before $v$, was *uninstalled* from the system (Section 2.2).

**Lemma 4** *Let $v$ to be a view installed in the system and $\mathcal{G}^v$ its associated view generator. If $\mathcal{G}^v$ generates $seq : v_1 \rightarrow ... \rightarrow v_k \rightarrow w$, then only $w$ will be installed in the system.*

*Proof:* Given the sequence $seq : v_1 \rightarrow ... \rightarrow v_k \rightarrow w$ generated by $\mathcal{G}^v$, we prove by induction on $k$ that views $v_1$ to $v_k$ will not be installed. The base case ($k = 1$) is trivial. Servers in $v_1$ receive the state from $v$ (line 15) and do not install this view since the condition of line 21 is satisfied. Therefore, each $i \in v_1$ starts its view generator $\mathcal{G}_i^{v_1}$ by proposing the same sequence $seq_{v_1} : seq \setminus \{v_1\}$ (lines 22-23), which is eventually generated by $\mathcal{G}^{v_1}$ since all proposals are equal to $seq_{v_1}$. In the induction step, we have to show if the claim is valid for $k-1$, it will be valid for $k$. When processing $seq_{v_{k-1}} : v_k \rightarrow w$, each server in $v_k$ receives the state from $v_{k-1}$ and will not install $v_k$ since the condition of line 21 will be true. Consequently, views $v_1$ to $v_k$ will not be installed. Now, let us show that $w$ is installed. Since servers in $v_k$ receive the sequence $seq_{v_{k-1}}$, the view generators $\mathcal{G}^{v_k}$ will be started with a proposal $seq_{v_k} : seq_{v_{k-1}} \setminus \{v_k\} = \{w\}$, which eventually will be generated. During the processing of $seq_{v_k} = w$ the predicate of line 21 will be false and $w$ will be installed (line 25). Thus, only $w$ from $seq$ is installed in the system. □

**Lemma 5** *If a view $w$ is installed in the system, any previously installed view $w' \subset w$ was uninstalled and will not be installed anymore.*

*Proof:* Consider that $w$ is installed in the system. In this case, the view generators $\mathcal{G}^v$ associated with some view $v$ have defined a sequence $seq : v_1 \rightarrow ... \rightarrow v_k \rightarrow w$ for installing $w$ (Lemma 4). Consequently, all other views in $seq$ are uninstalled (line 18). Now, consider that a previous sequence $seq' : v_1 \rightarrow ... \rightarrow v_y \rightarrow w'$, $y < k$, generated by $\mathcal{G}^v$ led to the installation of another view $w'$. From the Weak Accuracy property of $\mathcal{G}^v$, $seq' \subseteq seq$ and $w' \in seq$. Consequently, the fact that $w' \in seq$ implies that $w'$ is uninstalled on the installation of $w$. From the Weak Accuracy of $\mathcal{G}^v$, any subsequent sequence $seq''$ generated by $\mathcal{G}^v$ wold have $seq \subseteq seq''$. Since the algorithm does not install outdated views (line 13), $w'$ will never be installed after $w$. □

**Lemma 6** *If a view $w$ is installed in the system, no other more up-to-date view $w'$ such that $w \subset w'$ was installed in the system.*

*Proof:* Assume there is a view $w'$, such that $w \subset w'$, and that this view is installed in the system. From Lemma 5, $w$ should not be installed, which implies a contradiction. Consequently, $w'$ is not installed in the system. □

**Lemma 7** *The views installed in the system form an unique sequence.*

*Proof:* Consider that a view $w_1$ is installed in the system. From Lemma 5, any previous outdated view is uninstalled from the system. From Lemma 6, no view more up-to-date than $w_1$ is installed. Now, consider that a reconfiguration makes the system reconfigure from $w_1$ to another view $w_2$. It means that $w_1 \subset w_2$ (line 13), that $w_1$ was uninstalled from the system (Lemma 5) and that no view more up-to-date than $w_2$ was installed (Lemma 6). The same argument can be applied inductively to any number of views, and thus it is possible to see that the views installed in the system form an unique sequence $w_1 \rightarrow w_2 \rightarrow ... \rightarrow w_{k-1} \rightarrow w_k$. □

**Lemma 8** *The reconfiguration procedure always terminate.*

*Proof:* Consider the system starts from an initial view $v_0$. The system reconfiguration starts by initializing view generators $\mathcal{G}^{v_0}$ (line 7), that will define some sequence for updating $v_0$ (lines 8-9). Lemma 3 states that the number of sequences generated by $\mathcal{G}^{v_0}$ is finite. The reconfiguration started for $v_0$ is going to terminate when the last view $w$ of the last sequence is installed. In the same way, a reconfiguration started by $\mathcal{G}^w$ makes the system reconfigure from $w$ to other updated view and finishes. However, by Assumption 4, there will be some view $v_{final}$ such that $\mathcal{G}^{v_{final}}$ will never be started since $\texttt{RECV} = \emptyset$ in all servers of $v_{final}$ (line 7). At this point the reconfiguration procedure terminates. □

## 6.2 FREESTORE Properties

The following lemmata and theorems prove that Algorithm 3 and the ABD adaptation for FREESTORE presented in Algorithms 4 and 5 satisfy the properties of Definition 1.

4. We consider $\mathcal{L}$ because it may generate more than just one sequence of views (Weak Accuracy). These proofs are also valid for Algorithm 3 with $\mathcal{P}$ since Strong Accuracy is a special case of Weak Accuracy (see Section 3).

#### 6.2.1 Liveness

Let us now show a sketch of proof that FREESTORE satisfies the liveness properties. Consider that $V(t)$ is the most up-to-date view installed in the system.

**Theorem 3** *(**Storage Liveness**): Every read/write operation executed by a correct client eventually completes.*

*Proof:* Consider a read/write operation $o$ executed by a correct client $c$. For each phase *ph* from $o$, $c$ sends a message to all servers in its current view $cv_c$ and waits for a quorum of responses (lines 5,14,22 and 32 – Algorithm 4; lines 1, 5 an 6 – Algorithm 5). In this scenario, we have three cases.

1) Case each server $i$ of this quorum replies with a view $v_i = cv_c$ (lines 7,15,24 and 33 – Algorithm 4), $c$ ends *ph* in $cv_c$ since $cv_c$ has at least a quorum of correct servers (Assumption 1). Consequently, $c$ starts the execution of the next phase for $o$, and so on, until it finishes the last phase for $o$ in $cv_c$, finally concluding $o$.
2) Case some server $i$ of this quorum replies with a view $v_i \neq cv_c$, then $c$ updates $cv_c$ if it discovers a most up-to-date view (lines 8,16,25 and 34 – Algorithm 4) and reinitiates *ph*. From Assumption 4 and Lemma 8, there is a time $t$ such that the system stops executing reconfigurations. Consequently, the last reconfiguration executed in the system lead to the installation of some up-to-date view $w$ ($V(t) = w$) (Lemma 8). Since servers send $w$ to $c$ that updates its $cv_c$ to $w$, we reach the previous Case 1.
3) Case $o$ is concurrent with a reconfiguration and servers blocked its execution at line 11 of Algorithm 3, then some up-to-date view is going to be installed (Lemma 8) and the servers unlock its execution at line 25 of Algorithm 3. Since $cv_c$ is not up-to-date we reach the previous Case 2.

Therefore, in the three possible cases $o$ eventually completes and the theorem follows. □

**Lemma 9** *Consider that a correct server $i$ executes a join/leave operation at time $t$, by sending an update $u$ to the servers in $V(t)$. Then, there will be a view $V(t')$, such that $V(t') \neq V(t)$ and $t' > t$, where $u \in V(t')$.*

*Proof:* Considering that a correct server $i$ send $u$ to the servers in $V(t)$ at $t$, two cases are possible.

1) Case all servers in $V(t)$ start their view generators associated with $V(t)$ proposing a sequence with a view that contains $u$ (line 7 – Algorithm 3). In this case, the fact that all proposals contains $u$ ensures that there will be a view $V(t')$, such that $t' > t$ and $V(t') \neq V(t)$, where $u \in V(t')$.
2) Case some servers in $V(t)$ start their view generators associated with $V(t)$, before the receipt of $u$, proposing a sequence with a view that does not contain $u$ (line 7 – Algorithm 3). In this case, considering a time $t' > t$ when a reconfiguration has installed a new view $V(t')$, two cases are possible: (a) $u \in V(t')$ since the proposals from the servers that received and proposed $u$ was computed in $V(t')$; or (b) $u \notin V(t')$ and $u$ was sent to all servers in $V(t')$ (lines 12 and 17 – Algorithm 3) that in the next reconfiguration will start its view generators with a proposal that contains $u$, reducing this situation to Case 1.

Therefore, in any case there will be a view $V(t')$, such that $V(t') \neq V(t)$ and $t' > t$, where $u \in V(t')$. Consequently, the lemma follows. □

**Theorem 4** *(**Reconfiguration – Join Liveness**): Eventually, the* enable operations *event occurs at every correct server that invoked a join operation.*

*Proof:* Consider that a correct server $i$ executes a *join* operation at time $t$, by sending an *update* $u = \langle +, i \rangle$ to the servers in $V(t)$. From Lemma 9, there will be a view $V(t')$, such that $V(t') \neq V(t)$ and $t' > t$, where $u \in V(t')$. Consequently, when $V(t')$ is installed at server $i$, the *enable operations* event occurs at this server, since the required condition ($i \in V(t')$) is satisfied for servers that are joining the system (lines 14 and 19 – Algorithm 3). □

**Theorem 5** *(**Reconfiguration – Leave Liveness**): Eventually, the* disable operations *event occurs at every correct server that invoked a leave operation.*

*Proof:* Consider that a correct server $i$ executes a *leave* operation at time $t$, by sending an *update* $u = \langle -, i \rangle$ to the servers in $V(t)$. From Lemma 9, there will be a view $V(t')$, such that $V(t') \neq V(t)$ and $t' > t$, where $u \in V(t')$. Consequently, when $V(t')$ is installed at server $i$, the *disable operations* event occurs at this server, since the required condition ($i \notin V(t')$) is satisfied for servers that are leaving the system (lines 14 and 27 – Algorithm 3). □

#### 6.2.2 Safety

Let us now show a sketch of proof that FREESTORE satisfies the safety properties. Firstly, we show that FREESTORE implements the safety properties of an atomic register [11], according to the following definition:

**Definition 3 (Atomic register)** *Consider two read operations $r_1$ and $r_2$ executed by correct clients. Consider that $r_1$ terminates before $r_2$ initiates. If $r_1$ reads a value $\alpha$ from register $\mathcal{R}$, then either $r_2$ reads $\alpha$ or $r_2$ reads a more up-to-date value than $\alpha$.*

Let $\mathcal{R}$ be the register, and $\omega(\beta)$ be the operation to write the value $\beta$ in $\mathcal{R}$ with an associated timestamp $ts(\beta)$. Consider that $V(t)$ is the most up-to-date view installed in the system.

**Lemma 10** *Each phase of a read/write operation always finishes in the most updated installed view.*

*Proof:* For a client $c$ to execute a phase $ph$ of a r/w operation, it is necessary that a quorum of correct servers be in the same view $v$ that $c$ holds (lines 7, 15, 24 and 33 – Algorithm 4; lines 1, 5 an 6 – Algorithm 5). Thus, we have to show that for any time $t$, $v = V(t)$. For the sake

of contradiction, consider that at time $t$, $c$ finishes $ph$ by receiving a quorum of replies from servers in $v = V(t') : V(t') \neq V(t) \wedge t' < t$. Without loss of generality, consider that $V(t')$ is the view immediately before $V(t)$ in the sequence of installed views (Lemma 7). Consequently, at least $V(t').q$ servers from $V(t')$ do not executed line 11 of Algorithm 3 for instalation of $V(t)$ since $V(t') \subset V(t)$ and they will stop to execute r/w operation until the update of their views to $V(t)$ (lines 11, 18 and 25 – Algorithm 3). However, to install $V(t)$ it is necesssary that at least $V(t').q$ servers from $V(t')$ execute line 12 of Algorithm 3 in order to update the state of servers in $V(t)$ (line 15 – Algorithm 3). Cleary, by quorum intersection properties, we reach a contradiction. Consequently, we have that $v = V(t)$ and the lemma follows. □

**Lemma 11** *Read/write operations always finish in the most updated installed view.*

*Proof:* By Lemma 10, the last phase of a read/write operation $o$ finishes in the most updated installed view. Consequently, $o$ also finishes in the most updated installed view. □

**Lemma 12** *Consider that $\alpha$ is the value from the last write $\omega(\alpha)$ completed in $V(t)$. Then, in the next reconfiguration $F$ from $V(t)$ to $V(t')$, $t' > t$, one of two cases may happen:*

*(1) If there is no concurrent write operation with $F$, then $\alpha$ is propagated to $V(t')$;*

*(2) If $F$ is concurrent with a write operation $\omega(\beta) : ts(\beta) > ts(\alpha)$, then either $\alpha$ or $\beta$ is propagated to $V(t')$.*

*Proof:* Consider that the last write $\omega(\alpha)$ completed in $V(t)$ has stored the value $\alpha$ in $\mathcal{R}$.

1) In case there is no concurrent write operation with $F$: at least $V(t).q$ servers store $\alpha$ and $ts(\alpha)$ is the greatest timestamp in the system (otherwise $\omega(\alpha)$ would not be the last). From the quorums intersection properties and since STATE-UPDATE messages are sent (lines 12 and 15 – Algorithm 3) and views installed in the system form an unique sequence (Lemma 7), the servers in $V(t')$ are going to receive $\langle \alpha, ts(\alpha) \rangle$ coming from $V(t)$ updating their register values (line 16 – Algorithm 3).
2) In case $F$ is concurrent with a write operation $\omega(\beta) : ts(\beta) > ts(\alpha)$: Three situations need to be considered:
   a) No servers in $V(t)$ updated their local state with $\langle \beta, ts(\beta) \rangle$ before the instalation of $V(t')$ (lines 2-4 – Algorithm 5) and we reduce this case to the Case 1.
   b) Some servers in $V(t)$ already store $\langle \beta, ts(\beta) \rangle$ sending this value to servers in $V(t')$, consequently it is possible that some servers from $V(t')$ receive only $\alpha$ while some others receive $\beta$ (line 16 – Algorithm 3). Anyway, $\omega(\beta)$ is going to be reinitialized in $V(t')$ (Lemmata 10 and 11) and all its servers will eventually store $\langle \beta, ts(\beta) \rangle$ in $\mathcal{R}$.
   c) All servers in $V(t)$ already store $\langle \beta, ts(\beta) \rangle$ sending this value to servers in $V(t')$, consequently all servers in $V(t')$ will store $\langle \beta, ts(\beta) \rangle$ in $\mathcal{R}$ (line 16 – Algorithm 3).

Since there is always a sequence of installed views (Lemma 7), the above cases are valid on the update of the next view $V(t')$ whenever a subsequent view $V(t'')$, $t'' > t'$, is installed and so on. Consequently, the lemma follows. □

**Lemma 13** *Consider that a read operation $r_1$ returns a value $\alpha_1$ at time $t_1^e$, which has an associated timestamp $ts_1$. A read operation $r_2$ started at time $t_2^s > t_1^e$ returns a value $\alpha_2$ associated with a timestamp $ts_2$ such that either (1) $ts_2 = ts_1$ and $\alpha_2 = \alpha_1$ or (2) $ts_2 \geq ts_1$ and $\alpha_2$ was written after $\alpha_1$.*

*Proof:* Consider that $r_1$ starts at time $t_1^s$ and that $r_2$ ends at time $t_2^e$, then $t_2^e > t_2^s > t_1^e > t_1^s$. Depending on the occurrence of a reconfiguration $F$ that installs a following view $V(t')$ in the unique sequence of intalled views (Lemma 7), we have the following possible, mutually exclusive, scenarios:

1) No concurrent reconfiguration. In this scenario we have no reconfiguration and thus $V(t_1^s) = V(t_1^e) = V(t_2^s) = V(t_2^e)$. Therefore, the behavior of the algorithm is the same of the basis protocol. Consequently, by the quorum intersection properties, the *write-back* phase from $r_1$ (lines 27-35 – Algorithm 4) ensures that $r_2$ reads $\langle \alpha_2, ts_2 \rangle$ from servers in $V(t_2^e)$ such that either (1) $ts_2 = ts_1$ and $\alpha_2 = \alpha_1$ or (2) $ts_2 \geq ts_1$ and $\alpha_2$ was written after $\alpha_1$.
2) Reconfiguration concurrent with $r_1$ ($t_1^s \leq t' < t_1^e$)[5]. Since $t_2^s > t_1^e > t'$, the reconfiguration installs $V(t')$ before both $r_2$ starts and $r_1$ ends. Consequently, from Lemmata 10 and 11, both $r_1$ and $r_2$ finish by reading from servers in $V(t') = V(t_1^e) = V(t_2^s) = V(t_2^e)$. For the same reasons of Case 1, $r_2$ reads $\langle \alpha_2, ts_2 \rangle$ from servers in $V(t_2^e)$ such that either (1) $ts_2 = ts_1$ and $\alpha_2 = \alpha_1$ or (2) $ts_2 \geq ts_1$ and $\alpha_2$ was written after $\alpha_1$.
3) Reconfiguration between $r_1$ and $r_2$ ($t_1^e < t' < t_2^s$). In this case, the reconfiguration installs $V(t')$ after $r_1$ had finished and before $r_2$ starts. By Lemma 12 and the *write-back* phase from $r_1$ (lines 27-35 – Algorithm 4), $\langle \alpha_1, ts_1 \rangle$ or a more up-to-date pair $\langle \alpha_x, ts_x \rangle$, with $ts_x > ts_1$, is propagated to $V(t')$. By Lemmata 10 and 11, $r_2$ ends by reading from servers in $V(t') = V(t_2^s) = V(t_2^e)$. By the quorum intersection properties, $r_2$ reads $\langle \alpha_2, ts_2 \rangle$ from servers in $V(t_2^e)$ such that either (1) $ts_2 = ts_1$ and $\alpha_2 = \alpha_1$ or (2) $ts_2 \geq ts_1$ and $\alpha_2$ was written after $\alpha_1$.
4) Reconfiguration concurrent with $r_2$ ($t_2^s \leq t' < t_2^e$). In this case, the reconfiguration installs $V(t')$ after $r_1$ had finished, after $r_2$ has been started and before $r_2$ has been terminated. By Lemma 12 and the *write-back* phase from $r_1$ (lines 27-35 – Algorithm 4), $\langle \alpha_1, ts_1 \rangle$ or a more up-to-date pair $\langle \alpha_x, ts_x \rangle$, with $ts_x > ts_1$, is propagated to $V(t')$. By Lemmata 10 and 11, $r_2$ ends by reading from servers in $V(t') = V(t_2^e)$. By the quorum intersection properties, $r_2$ reads $\langle \alpha_2, ts_2 \rangle$ from servers in $V(t_2^e)$ such that either (1) $ts_2 = ts_1$ and $\alpha_2 = \alpha_1$ or (2) $ts_2 \geq ts_1$ and $\alpha_2$ was written after $\alpha_1$.

5. Notice that it is impossible to have $t_1^e = t'$ or $t_2^e = t'$ since r/w operations are stopped while the reconfiguration installs $V(t')$ (line 11 – Algorithm 3) and, consequently, we can not have a view installation and a r/w operation finishing at the same time.

Since there is always an unique sequence of installed views (Lemma 7), the above cases are valid whenever a subsequent view $V(t'')$, $t'' > t'$, is installed and so on. This suffices to prove the lemma. □

**Theorem 6** *(**Storage Safety**): The read/write protocols satisfy the safety properties of an atomic read/write register.*

*Proof:* This proof follows directly from Lemma 13. □

**Theorem 7** ***Reconfiguration – Join Safety**: If a server $j$ installs a view $v$ such that $i \in v$, then server $i$ has invoked the $join$ operation or $i$ is member of the initial view.*

*Proof:* Consider that a server $j$ installs a view $v$ such that $i \in v$, then some server $z$ started its view generator associated with a previous view $v' \subset v$ proposing a view sequence containing $v$ (Lemma 4) such that the update $\langle +, i \rangle \in v$. In this case, $\langle +, i \rangle \in$ `RECV` at $z$ (line 7 – Algorithm 3). Consequently, $i$ sent a message to $z$ with the join request by executing the $join$ operation (lines 1-2 – Algorithm 3) or $i$ is member of the initial view. □

**Theorem 8** ***Reconfiguration – Leave Safety**: If a server $j$ installs a view $v$ such that $i \notin v \wedge (\exists v' : i \in v' \wedge v' \subset v)$, then server $i$ has invoked the $leave$ operation.*

*Proof:* Consider that a server $j$ installs a view $v$ such that $i \notin v$ and that $\exists v' : i \in v' \wedge v' \subset v$, then some server $z$ started its view generator associated with a previous view proposing a view sequence containing a view $v$ (Lemma 4) such that the update $\langle -, i \rangle \in v$. In this case, $\langle -, i \rangle \in$ `RECV` at $z$ (line 7 – Algorithm 3). Consequently, $i$ sent a message to $z$ with the leave request by executing the $leave$ operation (lines 3-4 – Algorithm 3). □

# 7 DISCUSSION

## 7.1 DynaStore vs. FREESTORE

This section discusses some differences between DynaStore [8] and FREESTORE. Although our focus is on comparing our approach with DynaStore, we comment the relationship between these works and the recent proposed SpSn [10] and SmartMerge [9].

### 7.1.1 Convergence Strategy

In DynaStore, the reconfiguration process generate a graph of views through which it is possible to identify a sequence of *established* views (see Figure 1, left). A view that is not established works as an *auxiliary view*, that must be accessed during r/w operations. For any established view $v$, the maximum number of views that can succeed it is $|v|$, i.e, each member of $v$ can generate a new view, moreover, new views representing the combinations of these views could also be generated. In contrast, reconfigurations in FREESTORE install only a single sequence of views (see Figure 1, right). In this case, different generated sequences of views are organized in an unique sequence of installed views. For each installed view $v$, our implementation of $\mathcal{L}$ bounds the number of generated view sequences to $|v| - v.q + 1$. SpSn and SmartMerge also install a sequence of ordered configurations (views).

### 7.1.2 Liveness

In DynaStore, a process executes a leave and halts the system. However, for any time $t$, there is a bound on the number of processes that can leave the system without compromising liveness. Let $F(t)$ be the set of processes that crashed until time $t$ and $J(t)$/$L(t)$ the set of pending joins/leaves at $t$, respectively. The liveness condition of DynaStore states that, fewer than $|V(t)|/2$ processes out of $V(t) \cup J(t)$ should be in $F(t) \cup L(t)$ [8]. In contrast, in FREESTORE a process that executes a leave should wait for the installation of the updated view (without itself). If this restriction is not respected, the leaving server is considered faulty. This approach specifies a bound on the number of leaves for FREESTORE: fewer than $|V(t)|/2$ processes out of $V(t)$ should be in $F(t) \cup L(t)$. SpSn does not specify a liveness condition and SmartMerge uses policies to prevent the generation of unsafe views, which in practice restricts the number of allowed leaves.

### 7.1.3 Normal Case Execution

In DynaStore, the reconfiguration process generate a graph of views through which it is possible to identify a sequence of *established* views. A view that is not established works as an *auxiliary view*, but must be accessed during r/w operations. For each view (established or not), DynaStore associates a *weak snapshot object* that is used to store updates. For any view $v$, its weak snapshot object $wso_v$ is supported by a set $S_v$ of $|v|$ static SWMR registers, i.e., one register for each member of $v$. During a r/w operation on $v$, $wso_v$ must be accessed twice to verify if some update was executed on $v$. Each access to $wso_v$ comprises two reads in each register of $S_v$. Thus, to execute a r/w on $v$, it is necessary a total of $4|v|$ (4 sequential, $|v|$ parallel) quorum accesses, with two communication steps each. SpSn and SmartMerge use a similar approach: before execute each r/w a set with $|v|$ SWMR registers must be accessed to check for updates. In contrast, the overhead introduced by FREESTORE on a r/w is the local verification if the client' view is equal to the current view of the servers.

### 7.1.4 R/W and Reconfiguration Concurrency

A r/w operation needs to traverse the graph of views generated by DynaStore to find a view where it is safe to be executed (the most up-to-date established view). During this processing, each edge of the graph must be accessed in order to verify if there is a more up-to-date view. For each view $v$, it is necessary to access its weak snapshot object, which requires $2|v|$ (2 sequential, $|v|$ parallel) quorum accesses. While the system is converging to some established view, the r/w operation does not terminate. Similarly, in SpSn and SmartMerge, a r/w concurrent with a reconfiguration also needs to access intermediary views to found the most up-to-date installed view. On the other hand, in FREESTORE, r/w operations are directed to the most up-to-date installed view $v$, without accessing any auxiliary view. However, after some sequence of views for updating $v$ is obtained, the r/w operations are stopped and can only terminate after the installation of the most up-to-date view of this sequence.

### 7.2 Performance of Read/Write Operations

Table 1 shows the number of communication steps demanded to execute a r/w operation with DynaStore (DS) [8], SmartMerge (SM) [9], SpSn [10] and FREESTORE (FS), for a process that handles an updated or outdated view. A r/w operation in FREESTORE requires at most a third of the number of communication steps required to execute it in DynaStore or SpSn and at most a half of the steps required in SmartMerge. More important, it matches the performance of static protocols in the absence of reconfigurations. Moreover, some steps of DynaStore, SmartMerge and SpSn are to access a set of static SWMR registers, while in FREESTORE each step acesses only a single register, demanding few messages.

TABLE 1
Communication steps of r/w operations.

| | *Updated View* | | | | *Outdated View* | | | |
|---|---|---|---|---|---|---|---|---|
| | DS | SM | SpSn | FS | DS | SM | SpSn | FS |
| *Read* | 12 | 8 | 14 | 2/4 | 19 | 16 | 22 | 4/6 |
| *Write* | 12 | 8 | 18 | 4 | 19 | 16 | 26 | 6 |

### 7.3 Consensus-free vs. Consensus-based Reconfiguration

Table 2 shows the number of communication steps required for processing a reconfiguration using consensus-free (DynaStore, SmartMerge, SpSn and FREESTORE with $\mathcal{L}$) and consensus-based (RAMBO and FREESTORE with $\mathcal{P}$) algorithms. We present the number of communication steps for the best case scenario – when all processes propose the same updates in a reconfiguration – and for the worst case – when each process proposes different updates. In order to simplify our analysis, we consider that no reconfiguration is started concurrently with the one we are analyzing. Similarly, we assume a synchronous execution of the Paxos protocol [3], which requires only three communication steps, for consensus-based algorithms.

TABLE 2
Communication steps for reconfiguration.

| *Consensus-free approaches* | *Best Case* | *Worst Case* |
|---|---|---|
| DynaStore [8] | 23 | $18\|v\| + 5$ |
| SmartMerge [9] | 11 | $6\|v\| + 5$ |
| SpSn [10] | 14 | $8\|v\| + 6$ |
| FREESTORE with $\mathcal{L}$ | 4 | $7\|v\| - 2v.q - 1$ |
| *Consensus-based approaches* | *Best Case* | *Worst Case* |
| RAMBO [5] | 7 | 7 |
| FREESTORE with $\mathcal{P}$ | 5 | 5 |

FREESTORE reconfiguration is significantly more efficient than previous consensus-based and consensus-free protocols. FREESTORE with $\mathcal{P}$ requires two less communication steps than RAMBO and, more importantly, FREESTORE with $\mathcal{L}$ outperforms other consensus-free approaches by almost an order of magnitude. In particular, FREESTORE with $\mathcal{L}$ presents the best performance among all considered reconfiguration protocols in the best case, which is expected to be the norm in practice. An open question is if this number constitutes a lower bound.

## 8 CONCLUSIONS

This paper presented a new approach to reconfigure fault-tolerant storage systems, which clarify the differences between relying or not to consensus for agreement in the next view to be installed. The main result is a protocol that is simpler and cheaper (in terms of communication steps for either r/w operations or reconfigurations) than already proposed solutions and that fully decouples the execution of r/w and reconfigurations. Other interesting result is that in the best case, consensus-free reconfigurations are cheaper than the consensus-based ones.

A final contribution of this work is the introduction of a new abstraction called view generator. We believe that exploring different instantiations of this abstraction and their properties is an important avenue for future work.

## REFERENCES


[1] D. Gifford, "Weighted voting for replicated data," in *7th ACM Symposium on Operating Systems Principles*, 1979.

[2] T. D. Chandra and S. Toueg, "Unreliable failure detectors for reliable distributed systems," *Journal of the ACM*, vol. 43, no. 2, pp. 225–267, 1996.

[3] L. Lamport, "The part-time parliament," *ACM Transactions Computer Systems*, vol. 16, no. 2, pp. 133–169, 1998.

[4] H. Attiya, A. Bar-Noy, and D. Dolev, "Sharing memory robustly in message-passing systems," *Journal of the ACM*, vol. 42, no. 1, pp. 124–142, 1995.

[5] S. Gilbert, N. Lynch, and A. Shvartsman, "Rambo: A robust, reconfigurable atomic memory service for dynamic networks," *Distributed Computing*, vol. 23, no. 4, 2010.

[6] J.-P. Martin and L. Alvisi, "A Framework for Dynamic Byzantine Storage," in *34th International Conference on Dependable Systems and Networks*, 2004.

[7] M. J. Fischer, N. A. Lynch, and M. S. Paterson, "Impossibility of distributed consensus with one faulty process," *Journal of the ACM*, vol. 32, no. 2, pp. 374–382, 1985.

[8] M. K. Aguilera, I. Keidar, D. Malkhi, and A. Shraer, "Dynamic atomic storage without consensus," *Journal of the ACM*, vol. 58, pp. 7:1–7:32, 2011.

[9] L. Jenl, R. Vitenberg, and H. Meling, "Smartmerge: A new approach to reconfiguration for atomic storage," in *Proc. of the 29th Int. Symposium on DIStributed Computing*, 2015.

[10] E. Gafni and D. Malkhi, "Elastic configuration maintenance via a parsimonious speculating snapshot solution," in *Proc. of the 29th Int. Symposium on DIStributed Computing*, 2015.

[11] L. Lamport, "On interprocess communication (part II)," *Distributed Computing*, vol. 1, no. 1, pp. 203–213, 1986.

[12] M. Herlihy, "Wait-free synchronization," *ACM Transactions on Programing Languages and Systems*, vol. 13, no. 1, 1991.

[13] P. Dutta, R. Guerraoui, R. R. Levy, and M. Vukolić, "Fast access to distributed atomic memory," *SIAM J. Comput.*, vol. 39, no. 8, pp. 3752–3783, Dec. 2010.

[14] R. Fan and N. Lynch, "Efficient replication of large data objects," in *Proceedings of the 17th International Symposium on DIStributed Computing (DISC)*, 2003.

[15] M. K. Aguilera, I. Keidar, D. Malkhi, J.-P. Martin, and A. Shraer, "Reconfiguring replicated atomic storage: A tutorial," *Bulletin of EATCS: The Distributed Computing Column*, 2010.

[16] C. Dwork, N. A. Lynch, and L. Stockmeyer, "Consensus in the presence of partial synchrony," *Journal of ACM*, vol. 35, no. 2, pp. 288–322, Apr. 1988.

[17] V. Hadzilacos and S. Toueg, "A modular approach to the specification and implementation of fault-tolerant broadcasts," Department of Computer Science, Cornell University, New York - USA, Tech. Rep., May 1994.